

FUZZY LOGIC METHOD AND APPARATUS FOR BATTERY STATE OF HEALTH DETERMINATION

Technical Field

5 **[0001]** This invention relates to fuzzy logic based methods and apparatus for providing an output which is a fuzzy logic function of a plurality of inputs. The invention has particular application to methods and apparatus for determining the states-of-health of electrochemical batteries. In this context, the invention provides systems which measure electrical
10 characteristics of a battery and use fuzzy logic rules to determine state-of-health (SoH) from the measured electrical characteristics.

Background

15 **[0002]** A wide variety of rechargeable batteries are used today (in this specification and the appended claims, the term battery includes individual electrochemical cells as well as batteries which combine a number of electrochemical cells). Rechargeable batteries power a wide range of devices including things as diverse as cellular telephones, forklift trucks, backup power supplies, electric carts, electronic test equipment and so on.
20

25 **[0003]** Rechargeable batteries have finite lives. The ability of a battery to hold a charge and to deliver its rated current degrades with use and with the passage of time. Eventually batteries must be replaced. Also, batteries of types which have certain chemistries require periodic servicing. For example, cycling certain nickel-based batteries without fully charging and discharging them can cause the formation of crystalline structures within the batteries. This phenomenon, known as "memory," decreases the ability of a battery to hold charge. A condition of low capacity caused by a memory effect can sometimes be cured by servicing the battery. Such servicing may
30 involve fully discharging the battery and then fully recharging it.

35 **[0004]** Many businesses use significant numbers of rechargeable batteries. In such cases it can be difficult to keep track of which batteries are

in need of replacing or servicing. There is a need for methods and apparatus for determining the states-of-health of rechargeable batteries. Such apparatus should be able to provide an assessment of the state-of-health of a battery quickly.

5

[0005] The state-of-health of a battery must be distinguished from the state-of-charge (SoC) and capacity of the same battery. Capacity is the maximum charge that the battery is capable of holding. The capacity of a properly functioning new battery is typically equal to or greater than the specified capacity published by its manufacturer. SoC is a measure of how much charge the battery is currently holding as compared to the battery's capacity. For example a battery having a capacity of 5.0 Ampere hours (Ah) which contains a charge of 2.5 Ah would have a SoC of 50%. State-of-health, by contrast, gives a more thorough picture of the overall condition of a battery by assessing how closely a battery meets its design specifications. For example, the battery referred to above would be considered to have a poor state-of-health if it had a specified capacity of 15 Ah. Assessing state-of-health however involves more than simply comparing a battery's current capacity with its specified capacity. For example, the battery referred to above would be considered to have a poor state-of-health if it had a specified capacity of 5.0 Ah, but had an undesirably high internal resistance.

[0006] One approach to measuring state-of-health quickly is to measure the internal resistance of a battery. In general, excessive internal resistance indicates that a battery is in poor shape. Internal resistance measurements take only a few seconds to complete and provide a reasonably accurate indication of a battery's condition, especially if a reference reading from a good battery is available for comparison. Unfortunately, internal resistance measurements alone can provide only a rough indication of a battery's state-of-health. Various battery conditions affect the internal resistance. For example, the internal resistance of a

battery typically decreases for some time after the battery has been fully charged. After a few hours the internal resistance settles to a lower value. Temperature also affects internal resistance values. Many batteries contain protection circuits that further distort measured internal resistance values.

5

[0007] The problem of measuring the state-of-health of batteries is compounded by the fact that there are a very large number of different types of batteries. Each is designed to behave in a different manner. Rechargeable batteries may have any of a variety of battery chemistry types. Some common rechargeable batteries include: nickel cadmium (NiCd), nickel-metal hydride (NiMH), sealed lead acid (SLA), lithium ion (Li-ion) and lithium polymer (Li-polymer). Descriptions of various types of batteries can be found in *Handbook of Batteries*, 2nd ed.1995, David Linden; McGraw-Hill, Inc. ISBN 070379211. Further, for each chemistry type, a variety of battery models are available from various manufacturers. Any system for measuring the state-of-health of a battery must take into account the specific type and model of battery in question.

10

15

20

[0008] Previous approaches to determining the health of batteries include smart batteries, battery monitoring systems which relate battery health to the internal impedance of the battery, and battery monitoring systems which relate battery health non-linearly to one or more electrochemical parameters of the battery.

25

30

[0009] Smart batteries are devices containing at least one battery integrated with a system for measuring the state-of-health or state-of-charge of that battery. Because the state-of-health measurement system is designed for use with a single specific battery the system does not need to be adaptable to different types of batteries. A smart battery typically includes provisions for monitoring electrical parameters such as voltage, current into the battery and current out of the battery. Smart batteries typically include a simple processor that applies predictive algorithms to estimate the

battery's current state-of-charge. An example of a smart battery is described in United States patent No. 6,072,299. This patent describes a smart battery with maintenance and testing functions which maintains information related to the battery's need for maintenance and monitors conditions that indicate when the battery has reached the end of its useful life and should be discarded.

[0010] Smart batteries have several limitations. One limitation is that a system for measuring the battery's state-of-health must be integrated with each battery. This reduces the cost effectiveness of smart batteries. Further, to keep smart batteries cost effective the capabilities of the monitoring circuitry are generally limited. Such simplified circuits do not always provide reliable results.

[0011] Various stand-alone battery monitoring systems have been proposed. Some such systems measure the value of at least one electrochemical parameter of a battery as that battery ages. The way that the value(s) of the electrochemical parameter (or parameters) change with time is then related to the state-of-health of the battery. For example, United States patent 5,977,750 discloses a system for determining the health of a battery by tracking the degradation of the full charge capacity of a battery over time. As a further example, international patent publication WO 00/19578 discloses a battery monitoring system which computes intersections between temporal extrapolations of curves of internal resistance versus temperature with a predefined internal resistance versus temperature limit curve, thereby determining an expected useful life and hence the current state-of-health.

[0012] A disadvantage of such systems is that they require that periodic measurements take place as the battery degrades with age. A history of measurements of the relevant electrochemical parameters of the battery must accumulate before the degradation in those electrochemical

parameters can be determined. Such battery monitoring systems cannot determine the state-of-health of a particular battery without having a history for that particular battery.

- 5 **[0013]** Some battery monitoring systems measure an internal impedance of the battery and attempt to relate battery health to the internal impedance. Such systems are described, for example, in Huet, *A review of impedance measurements for determination of the state-of-charge or state-of-health of secondary batteries*, Journal of Power Sources **70** (1998) 59-69.
- 10 To date, such approaches have either not been as reliable as is desirable or have been undesirably complex and difficult to automate.

- 15 **[0014]** Some battery monitoring systems attempt to obtain more reliable indications of state-of-health by measuring a number of parameters and using a non-linear partial least squares analysis to identify a best fit between the measured electrochemical parameters of a battery and some mathematically defined state-of-health. For example, such an approach is described in Byers et al., *J. Electrochemical Soc.* **126** (1979) p. 720. The non-linear partial least squares analysis approach is undesirably
- 20 computationally intensive. Further, using too few parameters provides insufficient data for proper correlation. Using too many parameters introduces noise. Either of these factors can severely hamper the reliability of a system based on performing least squares analysis.

- 25 **[0015]** In some prior battery monitoring systems, taking a measurement of a battery's state-of-health takes undesirably long. This is especially the case for systems which require that the battery under test be fully charged before a state-of-health measurement can be made.

- 30 **[0016]** There remains a need for a system for measuring the state-of-health of batteries which ameliorates disadvantages of the prior art. There is a particular need for such systems which can be adapted quickly to

accommodate new types of battery. There is also a particular need for such systems which can quickly and reasonably reliably determine a battery's state-of-health.

5 Summary of the Invention

[0017] This invention provides methods and apparatus for training fuzzy logic inference systems to produce outputs which indicate a characteristic of a test system in response to receiving a plurality of parameter values. The invention is described below using as an example of
10 such a system a battery testing system. In this case the characteristic may be a state of health of a test battery and the parameter values may be the values of parameters such as voltages, currents, and times measured during the application of a current waveform to a test battery.

15 [0018] One aspect of the invention provides a method comprising:
a) providing a prototype set of fuzzy logic membership functions, the prototype set comprising a plurality of membership functions corresponding to each of the input parameters;
b) obtaining parameter values from a calibration system for which the
20 characteristic has a known value;
c) for each of the parameter values obtained from the calibration system obtaining a system-specific set of fuzzy logic membership functions by scaling the corresponding plurality of membership functions; and,
d) using the system-specific set of membership functions to obtain
25 outputs indicative of the characteristic of test systems.

[0019] Where the system comprises an electrochemical battery and the characteristic comprises a state of health of the battery, providing the prototype set of fuzzy logic membership functions may comprise
30 determining a chemistry type of the battery and selecting from a plurality of prototype sets of fuzzy logic membership functions a prototype set of fuzzy logic membership functions which match the chemistry type of the battery.

[0020] Another aspect of the invention provides an automatic battery testing method which comprises training a fuzzy logic battery analyzer for assessing the state-of-health of batteries of a known model and chemistry type, the method comprising:

- a) providing a prototype set of prototype fuzzy logic membership functions which matches a chemistry type of a model of batteries to be tested;
- b) measuring a capacity of a battery of the model of batteries to be tested which has a known state of health;
- c) applying a current waveform to the battery and measuring a plurality of parameter values;
- d) for each of the parameter values obtaining a model-specific set of fuzzy logic membership functions by scaling a corresponding plurality of membership functions in the prototype set; and,
- e) using the model-specific set of membership functions in combination with a set of fuzzy logic rules to assess the states of health of one or more batteries to be tested.

[0021] Other aspects of the invention provide apparatus which may be used to perform methods according to the invention. The invention may be practised by providing a computer program product comprising computer-readable information which can be run on an automated testing apparatus, such as a programmable battery tester.

[0022] Further features and advantages of the invention are set out below.

Brief Description of Drawings

[0023] In drawings which illustrate non-limiting embodiments of the invention:

Figure 1 is a block diagram of an apparatus according to a preferred embodiment of the invention;

Figure 2 is a plot of current as a function of time for a possible current waveform for use in the invention;

5 Figure 3 is a flowchart illustrating a fuzzy logic method for determining the state-of-health of a battery under test;

Figure 4 is a flowchart illustrating a fuzzy logic method for to computing the state-of-health of a battery under test in accordance with its model type;

10 Figure 5A is a graphical representation of an example fuzzy logic membership function;

Figure 5B is a graphical representation of a set of three exemplary fuzzy logic membership functions for use in fuzzification of a resistance parameter;

15 Figure 6 is a flowchart illustrating a method for efficiently fuzzifying parameter values;

Figure 7 is a graphical representation of a set of three exemplary fuzzy logic membership functions for use in defuzzification;

20 Figure 8 is a flowchart illustrating a training method that uses a calibration battery to create and optimize a model-specific matrix for use with a new model of battery;

Figure 9 is a flowchart illustrating a portion of the training method of Figure 8 in more detail;

25 Figure 10 is a flow chart illustrating operation of the optimization portion of the training method of Figure 8;

Figures 11A and 11B are flow charts illustrating an optimization routine for one parameter;

Figure 12A is a graphical illustration of the processes of translating a fuzzy set;

30 Figure 12B is a graphical illustration of the processes of flexing a fuzzy set;

Figures 13A and 13B illustrate specific flex and translation operations that may occur in optimizing fuzzy sets for a parameter;

Figure 14 is a process for translating a fuzzy set for use in the invention; and,

5 Figure 15 is a process for flexing a fuzzy set for use in the invention.

Description

[0024] Throughout the following description, specific details are set forth in order to provide a more thorough understanding of the invention.

10 However, the invention may be practised without these particulars. In other instances, well known elements have not been shown or described in detail to avoid unnecessarily obscuring the invention. Accordingly, the specification and drawings are to be regarded in an illustrative, rather than a restrictive, sense.

15

[0025] The invention is described in the following description as being applied to assessing the state-of-health of batteries. Those skilled in the art will understand that the methods and apparatus disclosed herein could also be applied more generally to the determination of output values
20 which are fuzzy logic functions of a plurality of input values.

[0026] Figure 1 is a block diagram of a battery analyzer **100** according to a preferred embodiment of the invention. Battery analyzer **100** comprises a means for applying a current waveform to a battery under test, a
25 measuring means for measuring electrochemical parameters of the battery under test and an analysis means for deriving an estimate of a state-of-health of the battery from electrochemical parameter values measured by the measuring means. Battery analyzer **100** is connected to a battery under test **102**. Preferably battery analyzer **100** begins by measuring the state-of-charge (SoC) of battery **102**. State-of-health (SoH) testing is best performed
30 when the battery under test is neither fully charged nor fully discharged. A rudimentary SoC determination may comprise measuring the battery's open-

circuit voltage. If analyzer **100** determines that the SoC of battery **102** is not appropriate for SoH determination then analyzer **100** either charges or discharges battery **102** to correct this situation.

5 **[0027]** A time-varying current waveform is applied to battery **102** via a measurement device **204** operating under control of a controller **200**. An example of a possible current waveform **230** is depicted by its corresponding voltage waveform in Figure 2. The particular mechanism by which the current waveform is generated is not important to this invention.

10 Those skilled in the art are aware of a wide variety of suitable means for applying a current waveform to a battery under test including many such means which are described in the worldwide patent literature. In the illustrated embodiment, either a current source **210** or a load **214** can be connected to battery **102** by a switch **211**. A controller **200** generates and

15 sends control signals **202** which set switch **211**, current source **210** and load **214** as appropriate to apply the desired current waveform to battery **102**.

[0028] Controller **200** may comprise a general purpose computer, microcontroller, or the like and an associated program store **201** containing

20 software instructions which, when executed by controller **200** cause controller **200** to operate as described herein. In the alternative, controller **200** may comprise a field programmable gate array (FPGA), application specific integrated circuit, or the like configured to execute methods according to the invention.

25 **[0029]** As the current waveform is being delivered to battery **102**, control signals **202** cause measurement device **204** to take measurements of electrochemical parameters of battery **102** and to convey those electrochemical parameters **206** to controller **200** in the form of digitized

30 signals. Neither the particular construction of measurement device **204** nor the way in which electrochemical parameters **206** are conveyed to controller **200** are important to the invention. Those skilled in the art are aware of a

wide variety of suitable means for obtaining electrochemical parameters of a battery under test including many such means which are described in the worldwide patent literature.

5 **[0030]** Controller **200** computes and displays on a display **212** and/or stores in a data store **203** associated with controller **200**, an assessment of a SoH of battery **102**. The assessed SoH is derived from the values of the electrochemical parameters **206**. In a preferred embodiment, the electrochemical parameters **206** that are measured by battery analyzer **100**
10 include:

- internal resistance of battery **102**;
- voltage at output of battery **102**, after charging, with no load;
- voltage at output of battery **102**, after charging, with a load applied;
- open circuit voltage of battery **102**; and,
- 15 • times which elapse between various events which occur during the measurement procedure.

[0031] By way of example, measurement device **204** might apply the waveform of Figure 2 to battery **102** to obtain a set of electrochemical
20 parameters **206** which can be used to assess the SoH of battery **102**. Different protocols for obtaining electrochemical parameters may be used for different battery types and models. The protocol may be specified for each battery by a c-code which specifies a set of factors which permit an appropriate test sequence to be established. The c-code may specify, for
25 example, chemistry, voltage rating, capacity and other information about the model of battery being tested.

[0032] The example current waveform **230** comprises a resistance test portion **230A**, an charge portion **230B**, a rest portion **230C** and a discharge
30 portion **230D**. During resistance test portion **230A**, measurement device **204** measures the internal resistance R1 and open-circuit output voltage OCV1. R1 and OCV1 may be used to determine if battery **102** is suitable for further

testing. If the values of R1 and OCV1 are within acceptable ranges, the method of this invention proceeds to charge portion **230B**. During charge portion **230B**, battery **102** is charged with current supplied by current source **210** at a specified rate. For example, the charging rate might be $0.5 \times C$,
5 where C is the capacity of battery **102** (e.g. if battery **102** has a capacity of 2 Ah then, during charge portion **230B**, battery **102** is charged at a rate of 1 Ampere). Preferably the charging rate is specified for each model of battery to be tested. In any case, the charging rate used must not damage battery **102**. Charge portion **230B** preferably comprises a series of pulses.
10 If the battery under test is a nickel-based battery, then, during each pulse, battery **102** is discharged (e.g. connected to a load) for a short period, and charged for a short period. If the battery under test is a lead-based or lithium-based battery, then, during each pulse, battery **102** is rested for a short period, and charged for a short period. Battery **102** is allowed to rest
15 for a short period between periods when it is being charged or discharged.

[0033] During the last charge pulse of charge portion **230B**, measurement device **204** measures voltage V1, which is the voltage of battery **102** when it is not under load and voltage V2, which is the voltage
20 of battery **102** when it is under load. In the preferred embodiment of the invention, controller **200** bases its SoH determination, in part, on the difference, DV1, between V1 and V2. After rest portion **230C**, measurement device **204** measures the internal resistance R2 of battery **102**.

25 [0034] During discharge portion **230D**, battery **102** is discharged at a specified rate. Discharge portion **230D** preferably comprises a series of pulses. During each pulse, battery **102** is discharged for a short period. This may be accomplished by controlling switch **211** to briefly connect battery **102** to load **214**. Load **214** is preferably a resistive load. During each pulse,
30 battery **102** is allowed to rest for a short period. In the preferred embodiment of the invention, in each discharge pulse, measurement device **204** measures a voltage V3 of battery **102** when battery **102** is not under

load and a voltage V4 of battery **102** when battery **102** is under load. Most preferably, discharge portion **230D** continues for a minimum specified time (for example 15 seconds) before the difference between V3 and V4 is measured. Discharge portion **230D** continues until the difference between V3 and V4 reaches a threshold value. Preferably the threshold value is specified for each model of battery to be tested. For example, the threshold value may be 6 mV/cell for NiCd, NiMH and SLA battery types, and may be 18 mV/cell for Li-ion and Li-polymer battery types. The time T1 which has elapsed since the start of discharge portion **230D** is measured. At the end of discharge portion **230D**, the two voltages V3 and V4 are also measured. In the preferred embodiment of the invention, controller **200** bases its SoH determination, in part, on the difference, DV2, between V3 and V4.

[0035] After current waveform **230** has been applied to battery **102**, controller **200** processes the set of electrochemical parameters **206** for battery **102** which was collected by measurement device **204**. In the preferred embodiment, the set of electrochemical parameters **206** includes at least R2, V1, V2, DV1, V3, V4 DV2 and T1. Controller **200** uses these parameters to derive and convey to a user an assessment of the SoH of battery **102**. It will be appreciated that the basic methods and apparatus of this invention can be applied in situations where different sets of electrochemical parameters are measured. Further, the mechanism by which the electrochemical parameters are measured does not affect the operation of assessing the battery's SoH. All that is required is some measuring means for measuring a set of electrochemical parameters of a battery which is sufficient to characterize the SoH of the battery. The measuring means is preferably software configurable. In the embodiment illustrated in Figure 1, the measuring means comprises measurement device **204**, current source **210**, switch **211**, load **214**, and the software instructions which cause controller **200** to generate the control signals **202** which cause the set of electrochemical parameters to be measured.

[0036] Controller **200** implements a fuzzy logic system for obtaining an assessment of the SoH of battery **102**. Figure 3 illustrates a method **300** to obtain such an assessment according to the invention. Method **300** determines in block **301** the model of battery **102**. Model identification information may be obtained, for example, by way of a suitable user input **201**. If the model is not known to apparatus **100**, as determined at block **302**, then a training function is performed (block **304**). The training function is described below. If block **302** determines from the model identification information that the model is known to apparatus **100** then the model identification information is used to look up information regarding the battery's electrical characteristics (e.g. a configuration code) and a set of information which determines how the SoH is related to the measured electrochemical parameters. This set of information may be called a "model-specific matrix" appropriate to the model of battery **102**. The c-code and model-specific matrix may be stored in a data storage device, such as a RAM, ROM, flash RAM, hard disk, or the like.

[0037] In block **308** a test sequence which suits the c-code identified in block **306** is applied to the battery to obtain a set of measured electrochemical parameters **206**. Block **308** may be carried out as described above, for example. In block **310** an assessment of the battery's SoH is obtained by applying fuzzy logic rules to the set of measured parameters. In block **312** the assessment is displayed and/or stored for later use.

[0038] Block **310** is performed as illustrated in Figure 4. To derive an assessment of the SoH of battery **102**, the set of measured electrochemical parameters is fuzzified (block **402**). Fuzzification involves using the measured values of the electrochemical parameters **206** for battery **102** to determine antecedent membership values in a number of fuzzy sets.

[0039] Fuzzy sets are generalizations of conventional sets. In conventional set theory, a given item is either a member of a set or is not a member of the set. When using fuzzy logic an item may have fractional membership in a set as determined by a membership function for the set.

5 The principles of fuzzy logic are now well understood to those skilled in the art and will not be repeated here. Various textbooks on the topic of fuzzy logic have been published including Ross, *Fuzzy Logic With Engineering Applications*, McGraw Hill, 1995 and Kosko, *Fuzzy Engineering*, Prentice Hall, 1997 which are hereby incorporated herein by reference.

10

[0040] In the preferred embodiment of this invention, for each electrochemical parameter **206** in the set, the value of the electrochemical parameter **206** is used to compute the membership of that electrochemical parameter in a fuzzy set associated with each one of a number of possible outcomes. For example, where there are three outcomes, such as “excellent”, “good” and “poor”, fuzzification comprises computing the values of three membership functions for each parameter in the set. Each membership function yields a value in the range of 0 to 1 where 1 indicates full membership in the fuzzy set associated with the membership function and 0 indicates no membership at all. More or fewer membership functions may be associated with each parameter. In general, the number of membership functions associated with each parameter may be different.

15

20

[0041] The membership functions may take various forms. Gaussian fuzzy logic membership functions most realistically model natural behaviour, but necessitate extensive computation. In the preferred embodiment of this invention, each of the membership functions is a triangular function as shown in Figure 5A. Function **500** has a peak **502**, and left and right points of intersection with the abscissa **504L** and **504R**. Function **500** has left and right linear segments **506L** and **506R** each having a slope. Function **500** has a low region **508L** in which the parameter has a value between that of points **504L** and **502** and a high region **508R** between

25

30

point **502** and point **504R**. Such a triangular membership function can be completely specified by three numbers. For example, the function of Figure 5A can be completely specified by providing the parameter values corresponding to each of points **504L**, **504R** and **502**. Equivalently, the function can be specified by the slopes of segments **506L** and **506R** and the parameter value corresponding to peak **502**.

[0042] Figure 5B shows an example set of membership functions for internal resistance for a specific battery. It can be seen that the membership functions are not symmetrical and can overlap with one another. The dotted lines in Figure 5B demonstrate that a battery having an internal resistance of 500 m Ω has an antecedent membership value of 0 in the set named “R-good”, about 0.35 in the set named “R-excellent” and about 0.1 in the set named “R-poor”. These values indicate the degree to which the internal resistance of battery **102** belongs to each of these fuzzy sets.

[0043] As illustrated by the “R-poor” membership function of Figure 5B, a trapezoidal membership function is a special case of a triangular membership function wherein one of the left and right points **506L**, **506R**, is located on the abscissa at a parameter value of “infinity”.

[0044] The results of fuzzification block **402** may be represented by a working matrix as shown in Table I. The working matrix of Table I includes example antecedent membership values. Of course, the results of fuzzification may be held in any suitable data structure or structures.

TABLE I - Example Antecedent Membership Values			
Sets \rightarrow	Excellent	Good	Poor
Parameters \downarrow	(j=0)	(j=1)	(j=2)
R (i=0)	0.7	0.6	0.2
DV1 (i=1)	0.5	0.8	0.3

DV2	(i=2)	0.35	0.55	0.65
T	(i=3)	0.25	0.45	0.85

[0045] The particular membership functions to be used for a battery depend upon the type and model of battery being tested. Definitions for the fuzzy logic membership functions for use with a specific model of battery can be conveniently stored in a model-specific matrix as selected in block 306. TABLE II shows an example format for a model-specific matrix. The information in the model-specific matrix may be held in any suitable data structure or data structures.

TABLE II - Example model-specific matrix

Battery model		<i>NiMH Standard</i>		
Name of fuzzy set	Row	parameter value at point 504L	parameter value at point 502	parameter value at point 504R
R-excellent	2	300 mΩ	450 mΩ	600 mΩ
R-good	3	150 mΩ	270 mΩ	350 mΩ
R-poor	4	560 mΩ	700 mΩ	'infinity'
DV ₁ -excellent	5	65 mV	96 mV	120 mV
DV ₁ -good	6	10 mV	56 mV	75 mV
DV ₁ -poor	7	120 mV	144 mV	'infinity'
DV ₂ -excellent	8	75 mV	123 mV	160 mV
DV ₂ -good	9	30 mV	66 mV	80 mV
DV ₂ -poor	10	100 mV	187 mV	'infinity'
T-excellent	11	74 ms	80 ms	'infinity'
T-good	12	73 ms	78 ms	80 ms
T-poor	13	0 ms	73 ms	75 ms
excellent	14	80%	100%	110%

good	15	70%	75%	80%
poor	16	50%	60%	70%

[0046] It can be seen that the model-specific matrix of Table II stores information about the model of the battery associated with the model-specific matrix; information specifying membership functions corresponding to each parameter in the set of parameters used to assess battery SoH (in the example of Table II there are 12 membership functions, with three membership functions for each of four parameters); and information which specifies defuzzification membership functions as described below (in the example there are three defuzzification membership functions). Preferably the model-specific matrix also stores a charging rate to be applied during the charge portion **230B** and a threshold value and discharging rate to be applied during the discharge portion **230D** discussed above. Preferably the model-specific matrix also stores the values of calibration counters and an epoch counter as described below. Each of the membership functions is defined by values located in one of rows 2 to 13 of the matrix. It will be apparent to those skilled in the art that the information displayed in Tables I and II need not be arranged in the form of matrices, but any data structures suitable for use by controller **200**, for example, linked lists or the like, may be used.

[0047] Figure 6 illustrates one method **600** for efficiently obtaining the antecedent membership values for a number, N , of parameters each having a number, P , of corresponding fuzzy sets. After initialization in block **602**, the first parameter value and the first membership function for the first parameter value are provided in block **604**. In block **605**, the first parameter value is compared to the parameter values for points **504L** and **504R** from the model-specific matrix to determine whether the parameter value lies within region **508L** or **508R**. If not, then the antecedent membership value for the first membership function and first parameter is set to 1 in block

608C, indicating that the first parameter does not fall within the first membership function at all. An antecedent membership value of 1 will not affect later computations because, as explained below, the computations of block 404 use a minimum function unaffected by an antecedent membership value at its maximum value of 1. If the first parameter value falls within the range of the first membership function, then, in block 606, the first parameter value is compared to the parameter values for points 504L and 502 from the model-specific matrix to determine whether the parameter value lies within region 508L. If so, then in block 608A the antecedent membership value for the first membership function and first parameter is determined from the first parameter value, from the point 504L and from the slope 506L of the membership function in region 508L. If the first parameter value is not within region 508L, but in region 508R instead, then in block 608B the antecedent membership value for the first membership function and first parameter is determined from the first parameter value, from the point 504R and from the slope 506R of the membership function in region 508R.

[0048] In block 622, the resulting value is stored in a suitable data structure which could, for example, be an array having the general format of Table I. If block 624 determines that there are more membership functions to be evaluated for the current parameter value then the counter j is incremented in block 625 and the process is repeated until all membership functions for the current parameter value have been evaluated.

[0049] If block 624 determines that there are no more membership functions to be evaluated for the current parameter value then method 600 continues to block 626 where it is determined whether there are more parameter values to be processed. If so, the counter I is incremented and the counter j is reinitialized to 0 in block 627 and the process is repeated for the membership functions for the next parameter value. After all

membership functions for all parameter values have been evaluated then the fuzzification process ends.

5 [0050] Although not shown in figures 5A, 5B and 6, if an antecedent membership value is determined by the method of figure 6 to be below some minimum threshold, then a small, non-zero antecedent membership value such as 0.01 can be stored by block 622. A small, non-zero antecedent membership value avoids mathematical difficulties which could be caused by the value zero without introducing significant error into the
10 computations.

15 [0051] Those skilled in the art will appreciate that each of the antecedent membership values can be calculated independently of any of the other antecedent membership values. The antecedent membership values can therefore be calculated in any order or even simultaneously. The method 600 illustrated in Figure 6 is merely an illustration of one possible way to perform fuzzification.

20 [0052] As indicated by block 404, an assessment of the overall SoH of battery 102 is obtained by applying a set of fuzzy logic rules to the antecedent membership values resulting from the fuzzification performed in block 402. The fuzzy logic rules derive consequent membership values from the antecedent membership values obtained in block 402. In the preferred embodiment described herein, there is a fuzzy logic rule for every
25 possible combination of memberships of electrochemical parameters 206 in the set of outcomes. For example, where there are four parameters, each having potential for membership in three fuzzy sets there are $3^4=81$ rules.

30 [0053] In the preferred embodiment, each rule is of the form: "IF (Parameter₁ is in SET(1, k_1)) and (Parameter₂ is in SET(2, k_2)) ... and (PARAMETER_N is in set (N, k_N)) THEN Consequence_R". k_i are indices which can range over the P fuzzy sets corresponding to the outcomes

associated with each parameter. Each rule has a different set of k_i . The correspondence between the outcomes and the integer values of k_i are such that the outcomes are ordered according to their respective contributions to SoH. For example, where there are three fuzzy sets associated with each parameter (i.e. $P = 3$), k_i can each be an integer in the range of 0 to 2, with $k_i = 0$ corresponding to the outcome “excellent,” $k_i = 1$ corresponding to the outcome “good” and $k_i = 2$ corresponding to the outcome “poor.”

[0054] Where Parameter₁ is R, Parameter₂ is DV1, Parameter₃ is DV2, Parameter₄ is T and each parameter determines membership in sets associated with the outcomes “excellent,” “good” and “poor” then the rules may be expressed in the form:

RULE 1: If R is “excellent” and DV1 is “excellent” and DV2 is “excellent” and T is “excellent”, then C1.

RULE 2: If R is “excellent” and DV1 is “excellent” and DV2 is “excellent” and T is “good”, then C2.

RULE 3: If R is “excellent” and DV1 is “excellent” and DV2 is “excellent” and T is “poor”, then C3.

RULE 4: If R is “excellent” and DV1 is “excellent” and DV2 is “good” and T is “excellent”, then C4.

...

RULE 80: If R is “poor” and DV1 is “poor” and DV2 is “poor” and T is “good”, then C80.

RULE 81: If R is “poor” and DV1 is “poor” and DV2 is “poor” and T is “poor”, then C81.

[0055] C1 to C81 are consequent membership values. C1 to C81 are computed as the minimum value of the four antecedent membership values considered by the corresponding fuzzy logic rule. Thus, for example, RULE1 may be restated as follows:

RULE 1: C1 = the minimum of:

- the membership value of R in “excellent,”
- the membership value of DV1 in “excellent,”
- the membership value of DV2 in “excellent,” and,
- 5 - the membership value of T in “excellent.”

[0056] The consequent membership values are then defuzzified (block 406 in Figure 4) to yield an assessment of the SoH of battery 102. Defuzzification is performed by deriving from the consequent membership values a centroid of weights where each weight is computed as an area in a defuzzification membership function associated with the rule that generated the consequent membership value in question.

[0057] To associate each rule with a defuzzification membership function, the values of k_i may be used to assign each rule to a group and each group may be associated with a defuzzification membership function. In the preferred embodiment, the group to which a rule belongs is identified by summing the k_i which define the rule. In the example embodiment being described herein, there is one rule for which $\sum k_i = 0$, four rules for which $\sum k_i = 1$, ten rules for which $\sum k_i = 2$, sixteen rules for which $\sum k_i = 3$, nineteen rules for which $\sum k_i = 4$, sixteen rules for which $\sum k_i = 5$, ten rules for which $\sum k_i = 6$, four rules for which $\sum k_i = 7$ and one rule for which $\sum k_i = 8$.

25 [0058] In the preferred embodiment described above there are nine groups each having a distinct value $\sum k_i$. In general the number of groups will be the number of distinct values for $\sum k_i$. Expressed mathematically,

$$\text{Number of Groups} = N * (P - 1) + 1 \quad (1)$$

where N is the number of parameters and P is the number of corresponding fuzzy sets. For example, there will be 11 groups in the case where there are 5 parameters and 3 corresponding outcomes for each parameter.

5 **[0059]** Each group of rules may be associated with a definitional consequent membership value. Preferably, the definitional consequent membership values are monotonically ordered and fully cover a desired range of output values. Preferably definitional consequent membership values are equally spaced throughout the range, though a non-uniformly
10 spaced set may be used without departing from the spirit of the invention. Preferably the correspondence between the group of rules and the definitional consequent membership values are such that the contribution of each group of rules to SoH is consistent with the definitional consequent membership values.

15 **[0060]** In the preferred embodiment described herein, each group of rules corresponds to one of the definitional consequent membership values in a monotonically ordered set throughout the desired range of output values from 50 to 100%. This set of definitional consequent membership values
20 {50, 56.25, 62.5, 68.75, 75, 81.25, 87.5, 93.75, 100} is shown as definitional consequent membership values 711A through 711I in Figure 7.

[0061] In the preferred embodiment described herein, each definitional consequent membership value is associated with an outcome defined by
25 defuzzification membership functions. For example, rows 14, 15 and 16 of Table II, shown in Figure 7, define defuzzification membership functions corresponding to the outcomes "excellent," "good" and "poor." Preferably, the defuzzification membership functions do not overlap. The outcome associated with a group of rules is the outcome for which the definitional
30 consequent membership value falls between the upper and lower values of SoH defining the corresponding defuzzification membership function.

[0062] Table III shows example defuzzification associations between the groups of rules defined by $\sum k_i$, the definitional consequent membership values and the outcomes defined by the defuzzification membership functions of rows 14, 15 and 16 of Table II.

5

Table III - Example Defuzzification Associations		
$\sum k_i$	Definitional Consequent Membership Value (%)	Outcome
0	100	excellent
1	93.75	excellent
2	87.5	excellent
3	81.25	excellent
4	75	good
5	68.75	poor
6	62.5	poor
7	56.25	poor
8	50	poor

10

15

[0063] For example, in the preferred embodiment described herein, the k_i values corresponding to RULE 2 are 0, 0, 0 and 1 so that $\sum k_i = 1$. RULE 2 is therefore associated with the definitional consequent membership value 93.75. The definitional consequent membership value is associated with the outcome "excellent" because 93.75 falls between the values 80 and 110 of row 14 of Table II. (See also membership function 703 of Figure 7.) Thus, for defuzzification purposes, RULE 2 is associated with the outcome "excellent."

20

25

[0064] The consequent membership value that results from applying a rule defines a trapezoid in the membership function that corresponds to

the outcome associated with that rule. For example, consider a battery for which RULE1 produces a consequent membership value of 0.3. For RULE 1, $\sum k_i = 0$, and so RULE 1 belongs to the group associated with the definitional consequent membership value 100%. As indicated by the vertical dotted line in Figure 7, 100% corresponds to membership function **703** which corresponds to the outcome “excellent”. The consequent membership value of 0.3 defines a trapezoidal region **704** within function **703**. The area of region **704** can be considered to be a weight, associated with the consequent membership value of 0.3, to be assigned to the rule.

[0065] Defuzzification (block **406**) can be performed, for example, by the center of maximum (CoM) or center of area (CoA) methods. Where the CoA method is used, defuzzification involves identifying and computing the areas of the trapezoidal regions identified by each of the rules, summing the results for each of the groups of rules and then computing a centroid to yield an assessed SoH for battery **102**.

[0066] The centroid may be computed using the equation:

$$SoH = \frac{\sum_{groups} x_i w_i}{\sum_{groups} w_i} \quad (2)$$

where: x_i is the definitional consequent membership value associated with the i^{th} group of rules, and w_i is the total weight computed from the consequent membership values of the i^{th} group of rules.

[0067] The value SoH is an assessment of the SoH of battery **102**. In some cases it can be undesirably misleading to present to a user a seemingly “precise” value for SoH. In such cases it can be desirable to display to the user an assessment of SoH which describes the SoH of battery **102** in

words. For example, the battery might be described as being one of “excellent,” “good,” “fair” or “poor.”

5 [0068] Because the fuzzy logic membership functions, exemplified in Figure 5B, are specific to a given battery model, a set of fuzzy logic membership functions must be provided for each model of battery to be tested. The specific membership functions which will result in a reliable indication of the SoH for a battery is arrived at empirically through experimentation and analysis. Such sets may be stored as model-specific
10 matrices as described above with reference to Table II.

15 [0069] There will be many cases, however, where a battery must be tested which is not known to apparatus 100. To deal with this situation, this invention provides a novel adaptive system which can be used to generate model-specific matrices for batteries of types new to apparatus 100. The adaptive system includes one or more prototype matrices, each of which specifies a prototypical set of membership functions. Each prototype matrix illustrates the general degradation trend of a type of battery based upon the chemistry of the corresponding battery type. Apparatus 100 may include
20 prototype matrices for one or more different types of battery. The adaptive system performs a method in which it creates a model-specific matrix by modifying an appropriate prototype matrix.

25 [0070] The adaptive system may comprise software which runs on controller 200. Figure 8 illustrates a training method 800 according to the invention. Method 800 begins by selecting a battery type (e.g. NiCd) (block 802) and connecting a calibration battery of the selected type to apparatus 100 (block 804).

30 [0071] The calibration battery is then primed (block 806). Priming repeatedly cycles the calibration battery until its maximum capacity is reached. While priming is fairly time consuming it permits the capacity of

the calibration battery to be reliably assessed. The capacity of the battery can be determined by charging the battery and then timing the discharge of the battery. The actual SoH is assumed to be directly related to the battery's capacity. Preferably the first calibration battery is a brand new battery which
5 can be assumed to have a SoH of 100%. The SoH of other calibration batteries having different capacities can then be determined as follows:

$$SoH_{RM} = \frac{C_{RM}}{C_{BM}} \times 100\% \quad (3)$$

where C_{RM} is the reliably measured capacity of a calibration battery, C_{BM} ,
10 is the reliably measured capacity of a new calibration battery which is assumed to have a SoH of 100%. C_{BM} may be called a "benchmark capacity".

[0072] The reliably-measured state-of-health (SoH_{RM}) is then recorded in block **807**. In an alternative embodiment of the invention where the user
15 accurately knows the SoH of the calibration battery by some other means, computation of the SoH for the calibration battery may be foregone and the known SoH may be simply recorded at block **807** as input from the user. After recording the reliably-measured SoH_{RM} the calibration battery is partially discharged (block **808**) so that the battery is operating in a state
20 that is neither fully charged nor discharged. The battery is preferably allowed to rest, for example for ½ hour before continuing.

[0073] A current waveform, which is preferably the same as current waveform **230** described above, is then applied to the calibration battery,
25 the fuzzy membership functions defined in a prototype matrix associated with the battery type identified in block **802** are then used as described above to compute a SoH for the battery (block **810**). In block **812**, the computed SoH is compared to the reliably-measured SoH_{RM} recorded at block **807**. If these values agree to within a threshold (for example 2%) then

no optimization is necessary in respect of the current calibration battery (block **814**).

5 [0074] If block **814** determines that the computed SoH and reliably-measured SoH_{RM} values do not match then block **816**, which is described below, adjusts the prototype matrix in a way that attempts to minimize the difference between the reliably-measured SoH_{RM} recorded in block **807** and the SoH computed in block **810**. Different membership functions are adjusted depending upon the SoH_{RM} of the calibration battery. As indicated
10 by block **818** the preceding parts of method **800** may be repeated, preferably using calibration batteries that have different states-of-health. The modified prototype matrix is saved as a model-specific matrix in block **820**.

15 [0075] Preferably the model-specific matrix shown in Table II also stores the values of calibration counters and an epoch counter. Table V is an example of stored values of calibration counters and an epoch counter.

Table V - Example Calibration and Epoch Counters	
Calibration Category q	Calibration Counter c_q
1	1
2	0
3	0
4	1
25 5	0
6	0
7	0
8	1
Epoch Counter	3

[0076] Each calibration counter c_q records the number of training cycles (block 816) completed using a calibration battery of the associated calibration category q . The epoch counter is the sum of the calibration counters and tracks the total number of training cycles completed for the particular model of battery 102. The calibration counter c_q and the epoch counter are updated in block 1022 of Figure 10, as described below.

[0077] Figure 9 illustrates training method 816 in more detail. Method 816 begins at block 902 by assigning the calibration battery to a calibration category according to the reliably-measured SoH_{RM} recorded in block 807. Table IV shows a possible arrangement of calibration categories as is used in a currently preferred embodiment of the invention.

TABLE IV - Calibration Category and Scaling Parameters selected according to the value of SoH_{RM}

q, A and $B \rightarrow$	Calibration Category q	Scaling Parameter A	Scaling Parameter B
$SoH_{RM} (\%) \downarrow$			
$0 \leq SoH_{RM} \leq 59$	1	n/a	n/a
$59 < SoH_{RM} \leq 68.75$	2	59	68.75
$68.75 < SoH_{RM} < 74$	3	68.75	74
$74 \leq SoH_{RM} \leq 76$	4	74	76
$76 < SoH_{RM} \leq 81.25$	5	76	81.25
$81.25 < SoH_{RM} < 92$	6	81.25	92
$92 \leq SoH_{RM} \leq 100$	7	92	100
$100 < SoH_{RM}$	8	n/a	n/a

[0078] In block 904 a ratio z is obtained. The ratio z is used in scaling the points which define the membership functions. The ratio z may be calculated using the following equation:

$$z = \begin{cases} 1, & SoH_{RM} \leq 59 \\ \frac{B - SoH_{RM}}{B - A}, & 59 < SoH_{RM} \leq 100 \\ 1, & 100 < SoH_{RM} \end{cases} \quad (4)$$

where A and B are the scaling parameters obtained from columns 3 and 4 respectively of Table IV and SoH_{RM} is the reliably-measured SoH_{RM} recorded in block 807. In block 906, the ratio z obtained in block 904 is used to obtain a scaling factor for one or more membership functions of the prototype matrix.

[0079] How method 816 uses the reliably-measured SoH_{RM} for calibration depends upon whether or not the current calibration battery is the first calibration battery. Method 816 can determine whether the current calibration battery is the first calibration battery since the epoch counter for the first battery has a value of zero, as determined at step 908. If the calibration battery is a first calibration battery then, in step 912, the method derives, for the membership functions corresponding to each of the parameters, a scaling factor, α . All of the membership functions for the parameter may be scaled by the same scaling factor α_θ . The result is that a SoH computed for the calibration battery with the use of the membership functions matches approximately the reliably-measured SoH_{RM} for the calibration battery.

[0080] For each of the parameters, the method obtains a scaling factor by choosing a particular one of the points which define the membership functions for that parameter. Which one of the points is selected depends upon the calibration category of the calibration battery.

For example, Table V identifies points which may be used for each of the calibration categories of Table IV:

Table V - Scaling Points for Parameter <i>P</i>		
Calibration Category <i>q</i>	Membership Function	Point (see Figure 5A)
1	P-Poor	502
2	P-Poor	504L
3	P-Good	504L
4	P-Good	502
5	P-Good	504R
6	P-Excellent	504L
7	P-Excellent	502
8	P-Excellent	502

- 15 [0081] The scaling factor α_0 of block 912 may be calculated using the following equation:

$$\alpha_0 = \frac{[(1-z)x + zy]}{x} \quad (5)$$

- 20 where *z* is the ratio of equation (4), *x* is the current parameter value at the point identified in Table V for the current calibration battery, and *y* is the parameter value measured for the calibration battery.

- 25 [0082] In block 916, the three parameter values that define each membership function in the set of membership functions for each parameter (as shown for example in Table II) are each scaled according to a scaling function $S(l, m, r, \alpha) = (l', m', r')$. The scaling function $S(l, m, r, \alpha) = (l', m', r')$ for each set of membership functions may be written in matrix form as follows:

$$\begin{bmatrix} l' \\ m' \\ r' \end{bmatrix} = \alpha * \begin{bmatrix} l \\ m \\ r \end{bmatrix} \quad (6)$$

where l' , m' and r' are the three parameter values that define the triangular membership function (exemplified by points **504L**, **502** and **504R** respectively of Figure 5A) after scaling; l , m and r are the three parameter values that define the triangular membership function (exemplified by points **504L**, **502** and **504R** respectively of Figure 5A) to be scaled; and α is the scaling factor α_0 of equation (5). In general, a different scaling factor will apply to the membership functions associated with each of the different parameters.

[0083] For example, suppose a first calibration battery has a SoH_{RM} of 93%, the measured value for a parameter P is 50 and the set of membership functions in the prototype matrix which corresponds to that parameter is as set out in Table VI

Table VI - Membership Functions for Parameter P Before Scaling			
Function	l (504L)	m (502)	r (504R)
P-Excellent	10	20	30
P-Good	25	35	45
P-Poor	35	45	-

Intuitively, if 93% is excellent, the method should scale this set of membership functions so that the value “50” measured for parameter P in the calibration battery corresponds generally to the peak of the P-Excellent membership function. A battery having a SoH of 93% is in calibration category 7 (from Table IV). Table V indicates that for calibration category 7 the scaling point is point **502** of the P-Excellent function. Therefore x has the value 20. The value of z in this case is 0.875 from equation (4) and from equation (5) $\alpha_0 \approx 2.3$. Multiplying

each of the points which define this set of membership functions by 2.3 and rounding to integer values yields:

Table VII - Membership Functions for Parameter <i>P</i> After Scaling			
Function	<i>l</i>	<i>m</i>	<i>r</i>
P-Excellent	23	46	69
P-Good	48	81	104
P-Poor	81	104	-

The sets of membership functions for the other parameters are scaled similarly.

[0084] The method then updates the counters (block **918**) and proceeds to the optimization of block **920**, as described below.

[0085] If the current calibration battery is not the first calibration battery (i.e. step **908** detects a non-zero value in the epoch counter) then only one of the points which defines the membership functions which correspond to each parameter may need to be changed. This is because the points which define the membership functions have already been scaled. Which one of the plurality (nine in the example being described here) of points which define the three membership functions associated with each parameter is scaled depends upon the calibration category to which the current test battery has been assigned. The point to be scaled is the same point identified in Table VI. A new value for the point is obtained as a weighted average of its current value and the parameter value for the current calibration battery. The new value for the point obtained in block **914** may be calculated using the following equation:

$$x_{new} = \frac{c \times x_{old} + [(1 - z)x_{old} + zy]}{c + 1} \quad (7)$$

where x_{new} is the new value for the point, x_{old} is the existing value for the point, z is the ratio of equation (4); and c is the value of the calibration

counter of Table V for the calibration category q of the current test battery (as determined in block 902).

5 [0086] For example, if $y = 130$, $c = 0$, $z = 1$ and $x_{old} = 100$ then $x_{new} = 130$. If $y = 130$, $c = 5$, and $x_{old} = 100$ and $z = 1$ then $x_{new} = 105$.

10 [0087] After scaling has been performed, block 918 increments the calibration counter c_q corresponding to the calibration category selected in block 902. Block 918 also increments the epoch counter to maintain its value as the sum of the values of all c_q associated with the particular modified prototype matrix or model-specific matrix in question. After coarse calibration is completed then an optimization is performed as indicated generally by block 920.

15 [0088] As shown in Figure 10, block 920 first checks to see if further optimization is required. If block 1002 determines that there is already a good match between the reliably-measured SoH_{RM} of the calibration battery and the SoH computed using the existing scaled prototype matrix then further optimization is not required. Block 1002
20 may determine, for example, whether these values differ by more than a threshold amount (e.g. 2%) from one another. If not then the optimization routine terminates. If block 1002 determines that further optimization is required then an optimization is performed individually for each of the different parameters as indicated in blocks 1004 through 1010. Block
25 1020 stores the optimized values in a working copy of the prototype matrix. After optimization is completed and there are no more calibration batteries, block 820 of figure 8 stores the working copy of the prototype matrix as a model-specific matrix.

30 [0089] Figures 11A and 11B illustrate an optimization routine 1004 for one parameter, in this case internal resistance. Routine 1004 selectively translates and flexes the membership functions of Table II

corresponding to internal resistance. Figures 12A and 12B illustrate translation and flexing respectively. Routine **1004** begins at block **1102** by obtaining a difference between the reliably-measured SoH_{RM} for a calibration battery and the computed SoH obtained for the same battery using the current prototype matrix. Block **1104** determines whether this difference is positive or negative. If the difference is positive then the optimization blocks adjust the appropriate membership function(s) for that parameter so as to increase the computed SoH. If the difference is negative then the optimization blocks adjust the appropriate membership function(s) for that parameter so as to decrease the computed SoH.

[0090] Figure 11A illustrates the situation wherein the difference is positive. Blocks **1106**, **1108** and **1110** cooperate to select the one of the P membership functions (in this case $P=3$) associated with the current parameter for which the parameter value for the calibration battery has the greatest antecedent membership value. In blocks **1106**, **1108** and **1110**, R1 refers to the antecedent membership value of internal resistance in the membership function “R-excellent”; R2 refers to the antecedent membership value of internal resistance in the membership function “R-good”; and R3 refers to the antecedent membership value of internal resistance in the membership function “R-poor”. If block **1106** determines that the resistance of the calibration battery already has strong membership in the “R-excellent” set then nothing further needs to be done for the resistance parameter.

[0091] The bracketed pairs of numbers in Figure 11A indicate row and column coordinates of a value in Table II. The row labelled “R-excellent” is numbered 2. The three values which define the locations of the vertices of the “R-excellent” fuzzy set are located at positions (2,1) (2,2) and (2,3) in Table II. If block **1108** indicates that the resistance parameter most strongly belongs to the “R-good” set (which is defined in row 3 of Table II) then blocks **1120** and **1122** determine whether the

resistance parameter value for the calibration battery is in the low region **508L** or the high region **508R** of the range of resistance values which belong to the “R-good” set. If the resistance parameter value is in low region **508L** then block **1124** translates the “R-good” set in the negative direction as illustrated by the dashed line indicated by **1302** in Figure 13A. This tends to increase the SoH computed for the calibration battery.

[0092] If the resistance parameter value is in high region **508R** then block **1126** flexes the “R-good” and “R-excellent” sets as indicated by dashed lines **1304** and **1306**. The dashed line 1304 is an example of flexing positively wherein point (3,3) is moved along the abscissa in a positive direction. The dashed line 1306 is an example of flexing negatively wherein point (2,1) is moved along the abscissa in a negative direction. Simultaneously flexing the “R-good” and “R-excellent” sets has the overall effect of increasing the computed SoH.

[0093] If block **1110** indicates that the resistance parameter most strongly belongs to the “R-poor” set (which is defined in line 4 of Table II) then block **1130** determines whether the resistance parameter value for the calibration battery is in the low region **508L** of the range of resistance values which belong to the “R-poor” fuzzy set. If so then block **1132** flexes the “R-excellent” and “R-poor” sets as illustrated by the dashed lines **1308** and **1310** in Figure 13A. This tends to increase the SoH computed for the calibration battery.

[0094] Figure 11B illustrates the operation of optimization routine **1004** in case the computed SoH is greater than the reliably-measured SoH_{RM} . Blocks **1150** and **1152** cooperate to select the one of the P membership functions (in this case $P=3$) associated with the current parameter for which the parameter value for the calibration battery has the greatest antecedent membership value.

[0095] If the maximum antecedent membership value is for the “R-excellent” set then blocks **1154** and **1156** cooperate to determine whether the resistance value for the calibration battery lies in low region **508L** or high region **508R** of the “R-excellent” set. If the resistance value is in low region **508L** then block **1160** causes the “R-good” and “R-excellent” sets to be flexed as indicated by dashed lines **1320** and **1322** in Figure 13B. If the resistance value for the calibration battery is determined to be in high region **508R** then “R-excellent” and “R-poor” sets are flexed as indicated by dashed lines **1324** and **1326** of Figure 13B.

[0096] If block **1152** indicates that the measured resistance for the calibration battery has the greatest antecedent membership value in the “R-good” set then blocks **1170** and **1172** cooperate to determine whether the resistance value for the calibration battery lies in low region **508L** or high region **508R** of the “R-good” set. If the resistance value is in low region **508L** then block **1174** causes the “R-good” set to be translated in the positive direction as indicated by dashed line **1330** in Figure 13B. If the resistance value is in high region **508R** then block **1176** causes the “R-good” and “R-excellent” sets to be flexed as indicated respectively by lines **1320** and **1326**. Optimization for the remaining parameters proceeds in an analogous manner.

[0097] Figure 14 illustrates a process **1400** which may be used to translate a fuzzy set. Process **1400** begins by establishing a step size in block **1402**. Block **1402** may, for example, select a step size which is 5% of the parameter value corresponding to a point in the fuzzy membership function to be affected by the translation operation. In block **1404** a multiple of the step size, for example, 10 times the step size is added to each of the points to be affected by the translation. For example, a multiple of the step size may be added to points **504L**, **502** and **504R** when translating membership function **500** of Figure 5A.

[0098] Block **1406** initializes a counter which is used to limit the number of repetitions of a loop (comprising blocks **1408** through **1422**). In the loop, each point affected by the translation is translated in block **1408** by the step size. The point is prevented from having a negative value by block **1410**. Block **1412** computes a SoH using a prototype matrix which has been modified to include the translated points. This computed SoH is compared to the reliably-measured SoH_{RM} in block **1414**. Block **1414** may, for example, compute a difference between the computed SoH and the reliably-measured SoH_{RM} .

[0099] If block **1416** determines that the magnitude of the difference is less than a threshold value, E , then the values of the translated points are stored (block **1418**) and E is set equal to the difference (block **1420**). Block **1422** ends the loop after a predetermined number of iterations. After the loop terminates, block **1424** checks to ensure that the points are within valid ranges. If so then block **1426** modifies the prototype matrix to include the values of the translated points which were most recently stored in block **1418**.

[0100] Figure 15 illustrates a process **1500** which may be used to flex a fuzzy set. Process **1500** may simultaneously flex one, two or more membership functions. Process **1500** begins by establishing appropriate step sizes in block **1502**. Block **1502** may, for example, select for each point to be affected by the flex operation a step size which is 5% of the parameter value corresponding to that point. In block **1504** a multiple of the step size, for example, 10 times the step size is added to each of the points to be affected by the flexing. For example, a multiple of the step size may be added to point **504L** or point **504R** or both when translating membership function **500** of Figure 5A.

[0101] Block **1506** initializes a counter which is used to limit the number of repetitions of a loop (comprising blocks **1510** through **1524**).

In the loop, each point affected by the flexing is translated in the appropriate direction by the step size determined for that point in block **1510**. Block **1512** computes a SoH using a prototype matrix which has been modified to include the translated points. This computed SoH is compared to the reliably-measured SoH_{RM} in block **1514**. Block **1514** may, for example, compute a difference between the computed SoH and the reliably-measured SoH_{RM} .

[0102] If block **1516** determines that the magnitude of the difference is less than a threshold value, E , then the values of the translated points are stored (block **1520**) and E is set equal to the difference (block **1522**). Block **1524** ends the loop after a predetermined number of iterations. After the loop terminates, block **1526** checks to ensure that the points are within valid ranges. If so then block **1530** modifies the prototype matrix to include the values of the translated points which were most recently stored in block **1520**.

[0103] Preferred implementations of the invention comprise computers running software instructions which cause the computers to execute a method of the invention. The invention may also be provided in the form of a program product. The program product may comprise any medium which carries a set of computer-readable signals containing to instructions which, when run by a computer, cause the computer to execute a method of the invention. The program product may be in any of a wide variety of forms. The program product may comprise, for example, physical media such as magnetic data storage media including floppy diskettes, hard disk drives, optical data storage media including CD ROMs, DVDs, electronic data storage media including ROMs, flash RAM, or the like or transmission-type media such as digital or analog communication links.

[0104] Where a component (e.g. an assembly, device, circuit, layer etc.) is referred to above, unless otherwise indicated, reference to that component (including a reference to a "means") should be interpreted as a reference to any component which performs the function of the described component (i.e., that is functionally equivalent), including components which are not structurally equivalent to the disclosed structure which performs the function in the illustrated exemplary embodiments of the invention. Accordingly, the scope of the invention is to be construed in accordance with the substance defined by the following claims.

[0105] As will be apparent to those skilled in the art in the light of the foregoing disclosure, many alterations and modifications are possible in the practice of this invention without departing from the spirit or scope thereof. For example,

- The invention could be practised by measuring different parameters than those discussed herein;
- The invention could be applied to a fuzzy logic system for ascertaining SoH from the results of electrochemical impedance spectroscopy measurements; and,
- The invention could be applied using non-triangular membership functions.

Accordingly, the scope of the invention is to be construed in accordance with the substance defined by the following claims.